

MATCHING FIBER SKELETON JUNCTIONS IN PULP MICROGRAPHS VIA MULTI-FEATURE AGGREGATION

JI LIN¹, TOQEER AHMED¹, LI FENG³, PENGFEI YIN³ AND YIQI LIU^{✉,1,2}

¹School of Automation Science & Engineering, South China University of Technology, Guangzhou 510640, China, ²Centre of Papermaking and Pollution Control, South China University of Technology, Guangzhou 510640, China, ³Xi'an North Huian Chemical Industries Co., Ltd., Xi'an 710302, China
e-mail: 202420116226@mail.scut.edu.cn, toqeerbouk43@gmail.com, 3049442593@qq.com, 3547531916@qq.com, aulyq@scut.edu.cn

(Received July 9, 2026; revised October 1, 2026; accepted October 3, 2026)

ABSTRACT

In pulp and paper manufacturing, fiber morphology is critical for process control and decision-making support with the help of micrographs identification. Automation identification of fiber morphology is usually hindered by the difficulties in separating fiber trajectories at skeleton intersections. Existing heuristic rules, typically based on directional continuity or width similarity, often frustrate on collapsed and entangled fibers in industrial pulps, then leading to systematic errors for the detection. These errors degrade control actions for downstream tasks, such as refining adjustments and furnish blending. This paper proposes a data-driven method that extracts four local geometric features at each junction breakpoint: direction angle, local width, cross-region relation, and curvature. These features are aggregated using a Kolmogorov–Arnold Network (KAN). On a synthetic benchmark with ground truth (15 curves per image), the proposed method is better than heuristic baselines and achieves 0.978 in terms of F1. Ablation confirms that all four features contribute positively to fiber morphology identification. By applying to industrial pulp micrographs, our method produces fiber length and width distributions that match those of a commercial analyzer (FS5, Valmet). Without retraining, the proposed method transfers to six independent pulp samples. These include both bamboo and softwood across multiple refining intensities. High measurement stability and full statistical discrimination of the six conditions ($\eta^2 \geq 0.99$) confirm its suitability. This method serves as a reliable, image-based sensing front-end for data-driven refining control and process decision-making.

Keywords: Feature engineering, Fiber analysis, Fiber intersection matching, Microscopic image processing.

INTRODUCTION

The microscopic structure of fiber populations, including fiber length, diameter, curl, and their assembly into networks, can determine the essential properties of cellulose-based materials (D'Amore *et al.*, 2010; Kayali *et al.*, 2025; Lindström *et al.*, 2024; Paunonen *et al.*, 2026; Dias *et al.*, 2024; Pennells *et al.*, 2022; Verho *et al.*, 2025). In pulp and paper manufacturing, these shape and structure features are primary factors that influence ending product quality and serve as process control and decision-making support. Ensuring consistent product properties, therefore, requires dependable, real-time fiber measurements. However, current industry practices still heavily rely on offline laboratory analyzers, which prevent closed-loop optimization. An automated, reliable image-based measurement system could fill this gap by analyzing fiber morphology from micrographs in real time and at the resolution required for advanced process control. Despite the development of methodologies, existing image analysis pipelines have not been established as validated measurement tools: the morphological data they

provide lack traceability, uncertainty assessment, or formal validation in measurement science. These research gaps motivate the present study.

Conventional image-based fiber analysis pipelines typically deal with segmentation, skeletonization, and morphological measurement as separate steps (Arafat *et al.*, 2025; Baez *et al.*, 2026; Lux, 2013). With the wave of AI, segmentation approaches shift from global thresholding (Otsu, 1979) and adaptive methods (Bradley and Roth, 2007) to advanced deep learning models, such as U-Net variants (Ronneberger *et al.*, 2015; Alejo and Beltrán, 2024; Qamar *et al.*, 2024). Skeletonization, usually carried out through morphological thinning (Lee *et al.*, 1994) or medial axis transforms (Blum, 1967), often work together with distance-transform-based width estimation (Pourdeyhimi and Dent, 1999; Soille, 2003). Despite ongoing improvements in each stage, the entire pipeline faces a core measurement challenge: at fiber intersections, skeletonization creates topologically ambiguous junctions that can break apart individual fibers or incorrectly connect different fibers (Huang *et al.*, 2016; Verho *et al.*, 2025). The

resulting issues are not just visual imperfections, but systematic measurement biases—fiber lengths tend to be underestimated, network connectivity overestimated, and important morphological features like curl and kink density are distorted (Muñoz *et al.*, 2025; Sten *et al.*, 2024). Micrographs are two-dimensional projections of the three-dimensional fiber suspension. Therefore, fiber intersections not only represent planar overlaps but also reflect spatial crossings in three dimensions. In classical stereology, the estimation of three-dimensional structural parameters from two-dimensional sections or projections is a central theme (Underwood, 1970; Howard and Reed, 2005). From a stereological perspective, accurately resolving these two-dimensional intersections is a prerequisite for reconstructing the three-dimensional fiber architecture and for estimating key morphological parameters, such as length and orientation distributions. These biases often go unchecked, leading to final measurements that are unreliable and difficult to trace back to their true values, which undermines their usefulness in process control or quality assurance. Addressing this measurement problem requires a junction matching approach with quantifiable and validated accuracy against ground truth.

Existing methods for resolving fiber intersection ambiguity mainly depend on heuristic rules applied to a limited set of handcrafted geometric features. Simple techniques reconnect skeleton fragments based on spatial closeness or directional consistency (Gaiselmann *et al.*, 2013; Huang *et al.*, 2016; Lux, 2013), while D’Amore *et al.* (2010) used Delaunay triangulation to describe the overall network topology but did not recover per-fiber morphological distributions. Verho *et al.* (2025) chose continuations at junctions by minimizing the bending energy of B-spline fits yet reported consistent underestimation of average fiber length due to unresolved junction failures. In fungal network analysis, Sten *et al.* (2024) paired branch points using geodesic distance and branch-out vector angles, achieving an F1 score of 0.89. For textile fibers, Deng *et al.* (2024) matched crossed yarn fragments by slope similarity, and topology-aware methods like ToFiE (Togo *et al.*, 2026) apply persistent homology to preserve global structure in 3D but do not resolve individual fibers. A shared and critical limitation among these approaches is the lack of a quantitative benchmark for junction-level matching accuracy against ground truth: without such a benchmark, the measurement error caused by mismatching cannot be quantified, and the resulting morphological data remains untraceable.

Although machine learning has been widely adopted for fiber segmentation (Baez *et al.*, 2026;

Chirakitsakul *et al.*, 2025; Togo *et al.*, 2026) and even for pulp particle classification with engineered features (Lindström *et al.*, 2024), data-driven methods remain conspicuously absent from the specific task of fiber topology reconstruction. The closest analogue exists in neuron morphology reconstruction, where deep classifiers have been used to distinguish branch points from crossing points in 3D skeletons (Chen *et al.*, 2021; Guo *et al.*, 2021). These approaches, however, rely on 3D volumetric data, strong geometric priors, and large annotated training sets that are unavailable in 2D pulp micrographs. In a preliminary conference study, we proposed a synergistic morphology–topology framework for pulp fiber suspensions, with a heuristic matching module combining direction angle and average width (Lin *et al.*, 2026). While that work demonstrated end-to-end feasibility, its matching accuracy at dense intersections was not quantitatively validated, and the heuristic scoring formula could not adapt to the morphological diversity of industrial pulps. These observations motivate the present study: a learnable matching framework capable of aggregating multiple geometric features and subject to rigorous quantitative validation is needed to transform fiber intersection resolution from a heuristic post-processing step into a traceable measurement operation.

To fill this gap, the present work proposes a learnable fiber intersection matching framework. Four local geometric features—direction angle, local width, curvature, and cross-region geometric relation—are extracted at each candidate breakpoint and aggregated by a lightweight neural network, instantiated here as a Kolmogorov–Arnold Network (Liu *et al.*, 2024), to produce an optimal matching score. The approach targets each fiber in pulp suspensions rather than dense fiber networks. To validate this framework, a progressive validation strategy that begins with quantitative benchmarking on a synthetic dataset with known ground truth is established, where the matching operation achieves an F1 of 0.978 and thereby provides internal calibration. Moving from calibration to instrument-level comparison, the resulting fiber morphology measurements are evaluated against a commercial optical fiber analyzer (FS5, Valmet), and systematic bias is quantified through relative deviation analysis against the FS5 reference, with the majority of metrics falling within 5%. The validated measurements are further applied to industrial pulp samples spanning multiple species and refining levels, and formal statistical testing confirms complete discrimination between all sample pairs ($\eta^2 \geq 0.99$, all pairwise comparisons significant at $\alpha = 0.05$). The core contribution of this work is the introduction of a data-driven, learnable matching framework that transforms heuristic fiber intersection

resolution into a traceable measurement operation. This framework is complemented by a three-level validation strategy—synthetic ground-truth calibration, instrument-level bias quantification, and statistical process-discrimination assessment—that collectively demonstrates the credibility of image-derived fiber morphology. The experimental evidence confirms that the method achieves fiber length and width measurements in strong agreement with the industry gold standard, establishing the method as a reliable, image-based sensing front-end for data-driven refining and blending control.

MATERIALS AND METHODS

OVERVIEW OF THE ANALYSIS PIPELINE

The proposed framework operates as a multi-stage pipeline. As illustrated in Fig. 1, the pipeline comprises five sequential stages, of which the first three and the last follow our prior conference framework. The fourth stage—breakpoint matching and fiber reconstruction—represents the core methodological contribution of this paper, where the original heuristic matching module is replaced by a data-driven approach that aggregates multiple features through a Kolmogorov–Arnold Network (KAN) classifier (Liu *et al.*, 2024). A detailed description of each stage follows.

MULTI-SCALE ADAPTIVE PREPROCESSING AND FIBER SEGMENTATION

Given a raw grayscale micrograph, the first objective is to segment fiber foreground from the background and classify fibers by physical scale. This multi-scale approach accommodates the substantial morphological differences in real pulp samples: large fibers with high aspect ratios are suited to skeleton-based analysis, while small fibers and fragments are better characterized by direct contour measurement.

Preprocessing begins with background subtraction to correct for spatially non-uniform illumination. Let $I(x,y)$ be the raw image intensity and $B(x,y)$ a calibration background image acquired under identical illumination without a sample. The corrected image is computed as:

$$I_{\text{corr}}(x,y) = 255 - |B(x,y) - I(x,y)| \quad (1)$$

Two complementary binary masks are then generated in parallel. The first mask uses Otsu’s method (Otsu, 1979), which finds an optimal threshold t_{Otsu} by maximizing inter-class variance:

$$B_{\text{Otsu}}(x,y) = \begin{cases} 255, & I_{\text{corr}}(x,y) \geq t_{\text{Otsu}}, \\ 0, & \text{otherwise}, \end{cases} \quad (2)$$

The second mask uses adaptive local thresholding. The corrected image is smoothed with a 9×9 box filter to estimate the local background, and a threshold map is constructed by scaling the filtered image by 0.975, i.e., $T_{\text{local}}(x,y) = 0.975 * \bar{I}_{\text{corr}}(x,y)$, where $\bar{I}_{\text{corr}}$ denotes the box-filtered result. The adaptive binary mask is then:

$$B_{\text{local}}(x,y) = \begin{cases} 255, & (x,y) \geq T_{\text{local}}(x,y), \\ 0, & \text{otherwise}. \end{cases} \quad (3)$$

The initial fiber mask is obtained as the logical conjunction of the two binary results:

$$B_{\text{init}}(x,y) = B_{\text{Otsu}}(x,y) \wedge B_{\text{local}}(x,y) \quad (4)$$

Scale-aware filtering based on connected-component analysis follows. Let A_i denote the area of the i -th connected component in the refined binary mask. Two area thresholds classify fibers into three categories:

$$\text{Class}(i) = \begin{cases} \text{Noise}, & A_i \leq T_{\text{small}}, \\ S_{\text{small}}, & T_{\text{small}} < A_i \leq T_{\text{large}}, \\ S_{\text{large}}, & A_i > T_{\text{large}}. \end{cases} \quad (5)$$

Based on our experimental datasets, we set $T_{\text{small}} = 15$ pixels and $T_{\text{large}} = 100$ pixels throughout this study. Small fibers are processed through a streamlined contour-based pathway: for each connected component in S_{small} , the perimeter P and enclosed area A are extracted, and fiber length and diameter are approximated as $L = P/2$ and $D = 2A/P$, respectively.

SKELETONIZATION AND LOCAL WIDTH ESTIMATION

For large fibers, skeletonization and local width estimation are conducted jointly. The thinning algorithm of Lee *et al.* (1994) is applied to each large-fiber connected component, reducing it to a one-pixel-wide skeleton K that preserves the original topology. To suppress spurious branches caused by boundary noise and intersections, all skeleton branches shorter than three pixels are removed by pruning. The Euclidean distance transform $D(x,y)$ is then computed over the binary fiber mask, giving the distance from each pixel to the nearest background pixel.




Fig. 1. Overall workflow of the synergistic morphology–topology fiber analysis framework.

TOPOLOGY NODE DETECTION AND JUNCTION DECOMPOSITION

Skeleton points are classified into endpoints, branch points, and pass points based on their 8-neighborhood connectivity. Let $N_8(p)$ be the number of skeleton pixels in the 8-neighborhood of point p . The classification rule is:

$$p \in \begin{cases} \text{Endpoint,} & N_8(p) = 1, \\ \text{Pass point,} & N_8(p) = 2, \\ \text{Branch point,} & N_8(p) \geq 3. \end{cases} \quad (6)$$

Branch points are processed to decompose the skeleton into independent fragments. For each branch point p , a cross-core region is defined by adaptive morphological dilation: the dilation radius is set by the local fiber width estimate, and a circular dilation centered at p is applied. The intersection of the dilated region with the original binary fiber mask defines the cross-core, from which all skeleton pixels are removed. This severs the skeleton into independent fragments $\{F_k\}$.

BREAKPOINT FEATURE EXTRACTION

For each candidate breakpoint $p \in P_{\text{candidate}}$, four local features are extracted to characterize the local

fiber geometry at the junction boundary, capturing both intra-fiber continuity and inter-fiber discriminative cues for the subsequent matching stage.

Local direction angle θ_p . To estimate the local fiber direction at breakpoint p , tracing proceeds from p along its skeleton fragment F_k . Following the fragment via 8-connected neighbors, a sequence of n consecutive skeleton points $Q_p = \{q_1 = p, q_2, \dots, q_n\}$ is collected. The centroid $\bar{q}$ of this point set is computed, and the direction angle θ_p is taken as the four-quadrant arctangent of the vector from p to $\bar{q}$, modulo 360° :

$$\theta_p = \arctan2(\bar{q}_y - p_y, \bar{q}_x - p_x) \quad (7)$$

Local average width w_p . The width at breakpoint p is the average of the distance transform over the traced point set Q_p :

$$w_p = \frac{1}{|Q_p|} \sum_{q \in Q_p} D(q) \quad (8)$$

Local curvature κ_p . Curvature measures the local bending of the fiber segment adjacent to the breakpoint. Using the Menger curvature estimator, κ_p is computed from the start (q_1), middle (q_m , $m = \lfloor n/2 \rfloor$), and end (q_n) points of the traced sequence:

$$\kappa_p = \frac{4 \cdot \text{Area}(q_1, q_m, q_n)}{\|q_1 - q_m\| \cdot \|q_m - q_n\| \cdot \|q_n - q_1\|} \quad (9)$$

Cross-region geometric relation φ_p . While the preceding features describe the local geometry at the breakpoint itself, this feature captures information about the intersection region. Let C_g be the set of all skeleton pixels in cross-core region g . The region centroid c_g is computed, and φ_p is defined as the direction vector from breakpoint p to $\bar{c}_g$:

$$\varphi_p = \arctan2(p_y - \bar{c}_{gy}, p_x - \bar{c}_{gx}) + 180^\circ \quad (10)$$

All breakpoint features are stored per breakpoint, containing pixel coordinates, the four feature values, and a group identifier g indicating the cross-core region. Both angular features, θ_p and φ_p , are represented as scalar degrees in the range $[0^\circ, 360^\circ)$. A direction just below 360° is therefore numerically distant from one just above 0° although the two are geometrically nearly identical. We use this scalar representation throughout the present study; alternative circular representations, which are continuous across that boundary, are discussed in "Conclusion and Future Work". Breakpoints within the same group are candidates for pairwise matching.

INTELLIGENT BREAKPOINT MATCHING

The objective of this stage is to determine, within each cross-core region, which breakpoint pairs should be reconnected. This reduces to a binary edge classification problem on a graph constructed per group.

Graph Construction

For each group g with n_g breakpoints, we construct a complete undirected graph whose vertices are the breakpoints and edges represent candidate continuations. Each vertex v_i holds a feature vector $f_i = [\theta_i, w_i, \kappa_i, \varphi_i]^T$. Each candidate edge e_{ij} connecting v_i and v_j is represented by the concatenation of both endpoint vectors:

$$\begin{aligned} e_{ij} &= [f_i^T, f_j^T]^T \\ &= [\theta_i, w_i, \kappa_i, \varphi_i, \theta_j, w_j, \kappa_j, \varphi_j]^T \end{aligned} \quad (11)$$

The target label $y_{ij} \in \{0,1\}$ indicates whether p_i and p_j belong to the same fiber.

Heuristic Baseline Method

The heuristic method from our prior conference work serves as a baseline. It uses only two features—direction angle θ and width w —and defines an ad hoc matching score. For a candidate pair (p_1, p_2) , the angle score is:

$$E_{\text{angle}} = |\theta_{p_1} - \alpha_{p_1 \rightarrow p_2}| + |\theta_{p_2} - \alpha_{p_2 \rightarrow p_1}| \quad (12)$$

where $\alpha_{p \rightarrow q}$ is the direction angle from p to q . The width score is defined as:

$$E_{\text{width}} = \frac{|w_{p_1} - w_{p_2}|}{w_{p_1} + w_{p_2}} \quad (13)$$

Candidate pairs with $E_{\text{width}} > 0.5$ are first filtered out, and the remainder are sorted by ascending $(E_{\text{angle}}, E_{\text{width}})$ and selected via greedy non-overlapping assignment.

KAN-Based Learned Matching

To overcome the limitations of the heuristic approach, we replace the handcrafted scoring function with a Kolmogorov–Arnold Network (KAN) trained to predict the probability that a candidate edge represents a true fiber continuation (Liu *et al.*, 2024). Unlike MLPs, which apply fixed activation functions at nodes, KAN parameterize learnable univariate functions along network edges using B-spline basis functions. The overall workflow—from feature vector construction to KAN scoring and greedy matching—is illustrated in Fig. 2.

Network architecture. The KAN takes an 8-dimensional edge feature vector as input and outputs a scalar matching score. The architecture consists of two hidden layers with 16 units each, yielding an 8–16–16–1 configuration. Formally, let the input be $e_{ij} \in R^8$. The network computes:

$$s(e_{ij}) = \Phi_3 \circ \Phi_2 \circ \Phi_1(e_{ij}) \quad (14)$$

where $\Phi_1 \in R^{8 \times 16}$, $\Phi_2 \in R^{16 \times 16}$, and $\Phi_3 \in R^{16 \times 1}$ are KAN layers. Each layer Φ_i maps an input vector x to an output vector y via learnable univariate functions $\varphi_{l,m,n}$ parameterized as B-spline curves of degree k :

$$y_m = \sum_n \varphi_{l,m,n}(x_n) \quad (15)$$

where $\varphi_{l,m,n}$ is the B-spline function connecting the n -th input to the m -th output in layer ℓ . This architecture




Fig. 2. Workflow of the learned breakpoint matching pipeline.

adaptively learns the functional relationships among the four breakpoint features, capturing nonlinear feature interactions that fixed heuristic rules or linear combinations cannot represent.

Loss function. Training is formulated as a binary classification problem with imbalanced classes: in a group of n_g breakpoints, the maximum number of true positive edges is $\lfloor n_g/2 \rfloor$, compared to $C(n_g, 2)$ total candidate edges. To address this imbalance, we employ binary cross-entropy loss with a positive class weight γ inversely proportional to the positive class frequency:

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{|E|} \sum_{e_{ij} \in E} [\gamma \cdot y_{ij} \log \sigma(s_{ij}) + (1 - y_{ij}) \log(1 - \sigma(s_{ij}))] \quad (16)$$

where $\sigma(\cdot)$ is the sigmoid function, E denotes the set of all candidate edges across all groups, and $\gamma = N_{\text{neg}} / N_{\text{pos}}$ is the positive class weight computed from the training data.

In addition, we introduce a symmetry regularization term that penalizes directional inconsistency in edge scores. Since a fiber continuation between breakpoints p_i and p_j is inherently undirected, the matching score should ideally satisfy $s(e_{ij}) = s(e_{ji})$, where e_{ji} is the reverse edge formed by swapping the endpoint feature vectors. We define the symmetry loss as the mean absolute difference between forward and reverse scores:

$$\mathcal{L}_{\text{SYM}} = \frac{1}{|E|} \sum_{e_{ij} \in E} |s(e_{ij}) - s(e_{ji})| \quad (17)$$

The total training objective combines both terms:

$$\mathcal{L} = \mathcal{L}_{\text{BCE}} + \lambda \cdot \mathcal{L}_{\text{SYM}} \quad (18)$$

where $\lambda = 0.2$ balances the symmetry constraint against the primary classification objective.

Greedy inference with mutual exclusivity. At inference, the trained KAN assigns a score s_{ij} to each candidate edge. Edges with non-positive scores ($s_{ij} \leq 0$) are discarded. The remaining edges are sorted in descending order, and a greedy algorithm iteratively accepts the highest-scoring edge, removes both endpoints from further consideration, and proceeds to the next candidate. This guarantees mutual exclusivity with complexity $O(n_g^2 \log n_g)$ per group.

Algorithm 1 summarises the inference pipeline of the proposed fiber intersection matching method.

TRAINING AND INFERENCE PROTOCOL

The KAN is trained exclusively on synthetic fiber data with known ground truth identities. Synthetic fiber suspensions are generated by a potential-field-driven algorithm that produces random curvilinear trajectories within a bounded canvas, each assigned a random width (1–3 pixels) and rendered with anti-aliasing and morphological dilation. A point–curve mapping is recorded, providing ground truth for breakpoint matching—two breakpoints are labeled as a true match ($y_{ij} = 1$) if they belong to the same generated curve and the same cross-core group.

Synthetic fiber suspensions were generated and processed through the same feature extraction pipeline. The resulting breakpoint features and ground truth labels were split 8:1:1 at the image level into training, validation and test sets, with normalization parameters computed from the training set and applied to the validation and test sets. The validation set was used for model selection and early stopping. The test set was used only for the final evaluation reported in Table 2; as the test-set results agreed with the validation results

Algorithm 1. Fiber intersection matching and reconstruction.**Input:** skeleton fragments $\{F_k\}$, cross-core regions $\{G_g\}$, trained KAN classifier Φ **Output:** Reconstructed fibers with length, diameter and curl

```

1: for each crosscore region  $g$  in  $\{G_g\}$  do
    // Stage 1: Feature extraction
2:   for each breakpoint  $p$  at the interface of  $g$  and fragment  $F_k$  do
3:      $f_p \leftarrow \text{ExtractFeatures}(p, F_k, g) \triangleright f_p = (\theta, w, \kappa, \varphi)$ 
4:   end for
    // Stage 2: Candidate scoring
5:   Construct all pairwise edge features  $e_{ij} \leftarrow [f_i; f_j]$ 
6:   for each edge  $e_{ij}$  do
7:      $s_{ij} \leftarrow \Phi(e_{ij}) \triangleright$  KAN matching score
8:   end for
    // Stage 3: Greedy assignment
9:   Sort edges in descending order of  $s_{ij}$ ; mark all breakpoints as free
10:   $M_g \leftarrow \emptyset$ 
11:  for each edge  $(i, j)$  in sorted order do
12:    if breakpoint  $i$  and breakpoint  $j$  are both free then
13:       $M_g \leftarrow M_g \cup \{(i, j)\}$ ; mark  $i$  and  $j$  as matched
14:    end if
15:  end for
16: end for
    // Stage 4: Fiber reconstruction
17: Link skeleton fragments according to all local matches  $M = \bigcup_g M_g$ 
18: for each reconstructed fiber  $r$  do
19:    $l_r \leftarrow \text{PhysicalLength}(r)$ 
20:    $d_r \leftarrow \text{MeanDiameter}(r)$ 
21:    $c_r \leftarrow \text{CurlIndex}(r)$ 
22: end for
23: return all  $(l_r, d_r, c_r)$ 

```

to within 0.01 F1, peak validation F1 is used as the reported metric. The KAN was optimized using Adam with a linear learning rate warmup from 10^{-5} to 10^{-3} over the first 100 epochs, after which the learning rate was held constant at 10^{-3} and reduced by half after 50 epochs without improvement, with early stopping at 300 epochs.

During inference, the trained KAN is applied identically to both synthetic validation images and real industrial micrographs.

FIBER RECONSTRUCTION AND PARAMETER CALCULATION

After breakpoint matching, matched edge pairs and unmatched fragments are combined into complete fiber trajectories. Each connected component of the intersection-removed skeleton serves as a base segment, and matched breakpoint pairs link these segments into longer chains.

For each reconstructed fiber L_m , three morphological parameters are computed. The physical length is the sum of Euclidean lengths of all constituent skeleton edges, weighted by the calibration factor s ($\mu\text{m}/\text{pixel}$):

$$\text{Length}_m = s \cdot \sum \|q_k - q_{k+1}\|_2 \cdot \delta(q_k, q_{k+1}) \quad (19)$$

where $\delta(\cdot, \cdot)$ is 1.4 for diagonal and 1 for orthogonal adjacency. The average diameter is computed from the distance transform along the fiber skeleton:

$$\text{Diameter}_m = 2s \cdot \frac{1}{|L_m|} \sum_{q \in L_m} D(q) \quad (20)$$

The curl index, measuring fiber tortuosity, is defined as:

$$\text{Curl}_m = \frac{\text{Length}_m}{\text{ChordLength}_m} - 1 \quad (21)$$

where ChordLength_m is the Euclidean distance between its two endpoints.

At the image level, where each micrograph contains many fibers, we aggregate these into three summary statistics: arithmetic mean length (L1), length-weighted mean length (L2), and length-weighted mean diameter (W1). These correspond to the standard industrial metrics (FL1, FL2, FW1) provided by the FS5 Fiber

Analyzer (Valmet), enabling direct comparison with the commercial reference analyzer.

RESULTS

The proposed methodology was evaluated on two complementary datasets: a synthetic fiber suspension dataset and a real industrial pulp fiber dataset. This dual evaluation enables controlled assessment with known ground truth alongside practical validation against an industry-standard device.

Synthetic dataset. Synthetic fiber suspensions were algorithmically generated. The number of curves per image was varied to produce a range of image complexities. As illustrated in Fig. 3, 15 curve-density levels were used, ranging from 10 to 20 curves (in steps of 1) and including 25, 30, 35, and 40 curves per image; these levels describe the complexity of the synthetic images and are not a statement about the suspension concentration. Each level was used to generate an independent dataset of 500 images, giving 7500 synthetic images in total, generated with a fixed random seed of 42. Each image is a 512×512 canvas with a 20-pixel margin. Curves are grown by a potential-field-driven random walk with a per-step angular perturbation of $\pm\pi/30$ and a centring drift of strength 0.1 decaying as $\exp(-d_{min}/10)$; each curve is assigned a length of 100–300 steps, a uniformly drawn start position and initial orientation, and a width drawn uniformly from 1.00 to 3.00 pixels, rendered by disc dilation. For every image, the pixel-to-curve mapping was automatically recorded to provide ground truth for breakpoint matching: two breakpoints are labeled as belonging to the same fiber if they originate from the same curve and the same intersection region.

Real industrial dataset. The real dataset consists of six sample groups from different pulp suspensions (Samples 1–6). Images were acquired with a Basler ace aca2440-75um camera fitted with a $0.5\times$ lens, at 2464×2056 pixels covering a physical field of $16.2 \text{ mm} \times 13.5 \text{ mm}$ (calibration factor $6.57 \mu\text{m}/\text{pixel}$). Fibers were imaged in transmission with backlight illumination, the pulp suspension was pumped continuously through a 3 mm thick flow cell during acquisition and 8-bit grayscale images were recorded, with a short exposure time to eliminate motion blur arising from the flow. Because the fibers are confined to a 3 mm thick cell, only those near the focal plane contribute to the imaged skeleton, and with fiber diameters of $17\text{--}30 \mu\text{m}$ an out-of-plane inclination shortens the projected length by only 0.4% at 5° and 1.5% at 10° . We therefore treat the fibers as lying essentially in the image plane; this is supported by the agreement of the measured

lengths with the FS5 reference, which is unaffected by projection effects. Each sample contains approximately 400–660 grayscale images (Table 4). As shown in Fig. 4 and summarized in Table 1, the six groups exhibit diverse fiber intersection densities, and average diameters.

Experimental environment. All experiments were conducted on a personal workstation equipped with an AMD Ryzen 9 7945HX processor (2.50 GHz), 32 GB of RAM, and an NVIDIA GeForce RTX 4060 GPU with 8 GB of dedicated memory. The method was implemented in Python, leveraging OpenCV for image processing operations, scikit-image for skeletonization and morphological analysis, NumPy for numerical computation, PyTorch for the implementation and training of the KAN classifier, and Pandas for data management and statistical aggregation.

Table 1. *Practical Application Datasets*

Sample No.	Fiber Characteristics
1	Low intersection, mean diameter $\approx 24 \mu\text{m}$
2	Moderate intersections, mean diameter $\approx 24 \mu\text{m}$
3	Complex intersections, longer fibers, coarse diameter $\approx 30 \mu\text{m}$
4	Dense intersections, shorter fibers, thin fibers ($\approx 17 \mu\text{m}$)
5	Moderate intersections, sparse fibers, mean diameter $\approx 21 \mu\text{m}$
6	Complex intersections, mean diameter $\approx 24 \mu\text{m}$

Naming conventions. Throughout this section, the heuristic matching method from our prior conference work is referred to as the Heuristic Matching method (HM), and the proposed data-driven matching method as the Learned Matching method (LM).

COMPARATIVE EXPERIMENT

To evaluate matching performance under controlled, scalable complexity, we tested all methods on synthetic fiber suspensions at 15 curve-density levels (10–40 curves per image). The KAN-based learned matching (LM) was compared against the heuristic baseline (HM) and four representative classifiers (MLP, RF, XGBoost, SVM) on identical train/validation/test splits; the primary metric is peak validation F1 score. All five learned classifiers were evaluated under comparable conditions: none was subjected to an extensive hyperparameter search, and all were run with representative standard settings on the identical candidate-edge set, ground-truth labels, eight-dimensional edge features and greedy mutual-exclusivity matching rule. Because class imbalance is




Fig. 3. Synthetic fiber suspensions with varying numbers of fibers.




Fig. 4. Example micrographs from the six real industrial samples.

the factor most likely to distort such a comparison, it was corrected consistently in every case: a positive-class weight in the KAN and MLP losses, a balanced class weight for SVM and random forest, and a positive-class scaling factor for XGBoost. The remaining settings were KAN 8–16–16–1 with B-spline degree 3 (Adam, learning rate $1e-3$ with warm-up and halving), MLP two hidden layers of 32 units with ReLU (Adam, learning rate $1e-3$), SVM with RBF kernel and standardized features, random forest with 200 trees of maximum depth 10, and XGBoost with 200 trees of maximum depth 6 at learning rate 0.1.

Table 2. Performance and inference speed of six methods on the synthetic fiber suspension dataset. F1 scores are averaged over 15 curve-density levels (10–40 curves per image). Inference speed is reported as seconds per 1000 breakpoints.

Method	Mean F1	Std	Speed (s/ 10^3 bp)
HM (heuristic)	0.8380	0.0201	0.0046
SVM (RBF)	0.9360	0.0079	1.5167
Random Forest	0.9410	0.0175	1.9995
MLP	0.9448	0.0168	0.0166
XGBoost	0.9580	0.0086	1.0521
KAN (LM, ours)	0.9656	0.0086	0.0672

Fig. 5 shows the 15 LM validation curves together with their mean (red) and $\pm 1\sigma$ band (pink). Training converges rapidly—mean F1 rises sharply within the first ~ 50 epochs and stabilizes by epoch 200—with a narrow standard deviation, indicating reproducible learning dynamics across curve densities. The inset and Table 2 compare peak F1 across methods. LM attains the highest mean F1 (0.9656) of the six methods and consistently exceeds the heuristic HM (0.8380) by a wide margin that grows with fiber density. All learned classifiers outperform HM, and the differences among them are small relative to the run-to-run variation: the second-ranked XGBoost is separated from the KAN by only 0.0076, comparable to their common standard deviation of 0.0086, so we do not treat the ranking among the learned classifiers as statistically established. MLP is the only learned baseline whose inference speed is comparable to that of the KAN (0.0166 s/ 10^3 bp, versus 0.0672 s/ 10^3 bp) and it is both less accurate and more variable. The distinguishing property of the KAN aggregator is therefore efficiency: among the learned classifiers it is the fastest that also ranks at the top on accuracy, and its inference is $15.7\times$ faster than XGBoost and $22.6\times$ faster than SVM. This arises from the compact 8-16-16-1 architecture, which processes all candidate edges in a junction simultaneously through dense batched computation,

whereas tree-based models evaluate edges sequentially through conditional branches.

The qualitative benefit of learned matching is evident in Fig. 6. Across three industrial micrographs and one synthetic suspension, LM reconstructs fiber trajectories that visibly match the manual ground truth, whereas HM frequently breaks fibers at dense crossings or merges adjacent ones. This visual improvement directly translates into the more reliable morphological measurements reported in the following sections.

FEATURE ABLATION STUDY

To disentangle the contributions of the four geometric features, we trained the KAN on six feature subsets: (A) all four features $\{\theta, w, \kappa, \varphi\}$; (B) w, θ only (matching HM’s information); (C) w, κ, φ ; (D) θ, κ, φ ; (E) θ, w, φ ; and (F) θ, w, κ . All configurations were trained for 600 epochs on a fixed synthetic dataset at a single fiber density (15 curves per image) under identical settings; their peak validation F1 scores are reported in Table 3, and the training curves appear in Fig. 7.

Table 3. Ablation configurations and F1 scores.

Config.	Features	F1 (synth)	F1 (real)
A	$\theta, w, \kappa, \varphi$	0.978	0.9382
B	θ, w	0.934	0.9245
C	w, κ, φ	0.890	0.8893
D	θ, κ, φ	0.912	0.9041
E	θ, w, φ	0.958	0.9272
F	θ, w, κ	0.966	0.9303
HM	θ, w	0.842	0.9144

F1 (synth) is the peak validation F1 on the synthetic dataset (Config. A–F), and F1 (real) is the F1 obtained by the same models on manual annotations of real junctions from the industrial samples (170 junction regions, 910 candidate edges, 270 true matches). HM denotes the heuristic baseline, which is evaluated on the same real-data annotation for reference.

The full-feature model (Config. A) achieves an F1 of 0.978, confirming that all four features contribute positively. Removing the direction angle θ (Config. C, $F1 = 0.890$, $\Delta = 0.088$) causes the largest performance drop, making it the single most discriminative cue—directional continuity remains the primary matching signal. Width consistency is the second most informative feature: omitting w (Config. D, $F1 = 0.912$, $\Delta = 0.066$) substantially degrades accuracy, reflecting the physical expectation that a fiber maintains approximately constant width along its trajectory. Curvature κ and cross-region relation φ contribute smaller but still positive gains




Fig. 5. Training dynamics of the KAN-based LM and peak F1 comparison across six methods. (a) Individual LM validation curves (light blue) together with their mean (red) and ± 1 standard deviation band (pink) plotted over 600 training epochs. A grey dashed line marks 0.9. (b) Inset: maximum (peak) F1 scores of LM (KAN), HM, MLP, RF, XGBoost, and SVM at each curve density. Markers denote the observed peak F1.

($\Delta = 0.020$ and $\Delta = 0.012$, respectively); curvature helps reject sharply bent segments that are unlikely to represent smooth continuations, while the cross-region angle provides a coarser spatial constraint. To confirm that these contributions are not an artefact of the synthetic benchmark, we additionally annotated junctions manually in a subset of the industrial samples. One hundred and seventy junction regions with pronounced crossing morphology were selected, yielding 910 candidate edges formed by all pairwise combinations of the skeleton breakpoints within each region, of which 270 were identified as true fiber continuations. On this real-data ground truth, all reduced configurations fall below the full four-feature configuration (Table 3), so no single feature can be discarded without a measurable loss of matching accuracy.

Two further comparisons are instructive. First, Config. B, which receives only θ and w but processes them through the learned KAN aggregator, attains an F1 of 0.934—substantially higher than HM’s 0.842 on the same inputs—confirming that data-driven aggregation alone discovers superior feature interactions. Second, the jump from Config. B to the full model (0.934 $\rightarrow$ 0.978) demonstrates the incremental discriminative

power contributed by κ and ϕ , which the heuristic baseline does not use.

INDUSTRIAL VALIDATION

To assess real-world measurement validity, we compared the morphological output of LM and HM against the FS5 Fiber Analyzer on six industrial pulp samples (Table 1). For each sample, three aggregate metrics—arithmetic mean length (L_1), length-weighted mean length (L_2), and length-weighted mean diameter (W_1)—were computed and compared to the corresponding FS5 reference values via relative deviation (Table 4).

LM achieves closer agreement with FS5 in 14 of the 18 sample-metric combinations, with relative deviations consistently below 7% and a maximum of 6.8% (Sample 4, L_1); the mean deviation across all combinations is 2.13%. HM, by contrast, reaches 12.9% in the worst case (Sample 2, L_2) and averages 4.27%. The four instances where HM appears closer are confined to L_1 or L_2 at differences below 2.5 percentage points and no systematic pattern. On W_1 , LM maintains deviations of 0.1–5.6% whereas HM spans 1.3–9.6%.




Fig. 6. Original micrographs, manual annotations, and reconstructed fiber populations from the heuristic matching method (HM) and the learned matching method (LM). The figure presents three industrial pulp samples and one synthetic fiber suspension.

LM achieves this agreement without any retraining on industrial images, the KAN having been trained exclusively on synthetic data. This indicates that the matching operation transfers from synthetic to real junctions; it does not by itself establish the absence of systematic measurement error. The close correspondence with the industry-standard FS5 positions the image-based pipeline as a credible candidate for near-real-time fiber quality monitoring.

CROSS-MATERIAL TRANSFERABILITY AND STATISTICAL RELIABILITY

To verify that LM transfers across pulp species and degrees of mechanical refining—beyond the samples

used for FS5 validation—and to subject its outputs to formal statistical inference, we processed six additional pulp suspensions without any retraining or parameter adjustment. The samples comprised three bamboo pulps at increasing levels of refining (Schopper–Riegler freeness: SR25, SR43, SR75) and three softwood pulps (SR20, SR30, SR73), each yielding over 1000 micrographs (Samples 7–12; Fig. 8).

Three levels of statistical analysis were applied.

(i) Bootstrap confidence intervals. Per-image L_1 and W_1 were bootstrapped (2000 resamples) to obtain 95% confidence intervals (Fig. 9). All intervals are remarkably narrow: standard errors are no greater than 0.003 mm for L_1 and no greater than 0.05 μm for




Fig. 7. *F1 curves for all ablation configurations.*

Table 4. *Practical Application Experimental Results*

Sample No.	Processed images	FL1 (mm)	LM L_1 (mm)	HM L_1 (mm)	FL2 (mm)	LM L_2 (mm)	HM L_2 (mm)	FW1 (μ m)	LM W_1 (μ m)	HM W_1 (μ m)
1	437	0.609	0.604	0.587	1.152	1.125	1.185	24.51	24.51	24.82
			(0.8%)	(3.6%)		(2.3%)	(2.9%)		(0.1%)	(1.3%)
2	572	0.609	0.618	0.610	1.161	1.180	1.311	24.32	24.47	24.72
			(1.5%)	(0.2%)		(1.6%)	(12.9%)		(0.6%)	(1.6%)
3	481	0.654	0.652	0.626	1.692	1.621	1.598	29.95	28.26	28.17
			(0.3%)	(4.3%)		(4.2%)	(5.6%)		(5.6%)	(5.9%)
4	608	0.602	0.643	0.640	0.817	0.853	0.875	17.37	17.73	19.03
			(6.8%)	(6.3%)		(4.4%)	(7.1%)		(2.1%)	(9.6%)
5	657	0.732	0.757	0.725	1.427	1.432	1.429	21.41	21.32	21.53
			(3.4%)	(1.0%)		(0.4%)	(0.1%)		(0.4%)	(0.6%)
6	614	0.630	0.625	0.617	1.197	1.16	1.238	23.68	23.67	21.71
			(0.8%)	(2.1%)		(3.1%)	(3.4%)		(0.1%)	(8.3%)

FL1, FL2, FW1 denote the arithmetic mean length, length-weighted mean length, and length-weighted mean diameter reported by the FS5 Fiber Analyzer (Valmet), and serve as the reference values. LML1/L2/W1 and HML1/L2/W1 are the corresponding metrics computed by the LM and HM methods, respectively, and are compared row-wise with FL1, FL2 and FW1. Values in parentheses are deviations from the FS5 reference.




Fig. 8. Representative micrographs of the six additional samples (Samples 7–12): (7-9) bamboo at SR25, SR43, and SR75. (10-12) softwood at SR20, SR30, and SR73.




Fig. 9. Bootstrap 95% confidence intervals of per-image mean fiber length (L_1) and length-weighted mean diameter (W_1) for the six test samples. Colored horizontal bars represent the percentile bootstrap CIs, black dots are sample means. (a) Bootstrap 95% CI for L_1 , (b) Bootstrap 95% CI for W_1 .




Fig. 10. Mini-batch scatter plot of L_1 versus W_1 for the six samples, with 95% confidence ellipses. Black dots mark the sample centroids. Letters indicate Games–Howell groupings ($\alpha = 0.05$): samples sharing a letter are not significantly different in L_1 .

W_1 (exact values in Table 5), confirming that the LM pipeline yields highly reproducible aggregate metrics despite substantial per-image variability.

(ii) Between-sample differentiation. Per-sample records were randomly partitioned into 20 mini-batches (50 images each), and Welch ANOVA was performed on the batch means. For L_1 , $F(5, 114) = 3663$, $p < 0.001$, $\eta^2 \approx 0.99$; for W_1 , $F(5, 114) = 27129$, $p < 0.001$, $\eta^2 \approx 0.99$. The six samples form well-separated clusters in the L_1 – W_1 plane (Fig. 10), and Games–Howell post-hoc tests confirm that every pairwise comparison in both metrics is statistically significant ($\alpha = 0.05$; distinct grouping letters a–f in Table 5). Effect sizes of $\eta^2 \geq 0.99$ indicate that virtually all observed variance is attributable to the known sample categories rather than to measurement noise. The method cleanly discriminates all six process conditions—spanning two fiber species and three degrees of mechanical refining—with every pairwise comparison reaching statistical significance.

(iii) Single-fiber distributions. The logarithmic length frequency distributions (Fig. 11a) reveal a systematic shift with increasing refining intensity (higher SR value): the proportion of very short fibers

(≤ 0.1 mm) rises while the long-fiber tail diminishes, producing monotonic increases in both skewness and the short-to-long fiber ratio R_N (Table 5). Width distributions for fibers longer than 0.1 mm (Fig. 11b) show that bamboo fibers maintain a stable width (~ 18 – 19 μm) across refining levels, whereas softwood fibers are distinctly wider and exhibit a modest width reduction with more intense refining. Kolmogorov–Smirnov tests confirm that every sample pair possesses significantly different length and width distributions ($p < 0.001$ for all pairwise comparisons).

The transfer to bamboo and softwood pulps, for which the synthetic training data have no explicit correspondence, indicates that the matching operation is driven mainly by local fiber geometry rather than by species-specific photometric or textural properties. This transferability is qualitative rather than proven, and it is relevant where furnish composition varies. Moreover, the ability to resolve the full fiber population distribution, including the heavy concentration of short fragments produced by refining, provides orthogonal process-sensitive descriptors (R_N , distribution skewness) that are unavailable from bulk optical analyzers such as the FS5 and could be exploited for real-time quality monitoring.




Fig. 11. Single-fiber size distributions. (a) Length frequency (log-scale vertical axis), inset magnifies the first bin (0–0.1 mm). (b) Width frequency.

DISCUSSION

The transition from a fixed heuristic formula to a learned aggregator changes how fiber intersection resolution is performed: instead of relying on two handcrafted features with predetermined weights, the network discovers feature interactions directly from data, and does so at a computational cost that is low enough for throughput-oriented use. Comparable heuristic strategies remain prevalent in fiber tracking across materials science, biology, and textile engineering (Deng *et al.*, 2024; Sten *et al.*, 2024; Verho *et al.*, 2025); the demonstrated efficacy of learned aggregation therefore suggests a broadly applicable pathway for improvement in these fields.

The ablation results further clarify the relative informativeness of the four geometric cues. Local

direction angle is the dominant matching signal, consistent with the well-established principle of directional continuity in fiber tracking. Local width consistency provides complementary information—a fiber tends to maintain approximately constant width along its trajectory—though its contribution, while substantial, is secondary. Curvature helps reject sharply bent junction fragments that are unlikely to represent smooth fiber continuations, while the cross-region geometric relation supplies a coarse spatial constraint that complements local direction information. More importantly, the comparison between the learned aggregator and the heuristic baseline, when both receive the same two features (θ and w), reveals a notable finding: data-driven aggregation alone, without introducing new features, already produces a substantial accuracy gain on the synthetic benchmark, where the

Table 5. Statistical summary of fiber morphology for six cross-material test samples.

Sample	L_1 (mm)	W_1 (μ m)	R_N^a	Length skewness	Width skewness ^b	Group ^c
	Mean $\pm$ SE [95% CI]	Mean $\pm$ SE [95% CI]				
BSR25	0.945 $\pm$ 0.003	19.02 $\pm$ 0.02	2.46	4.05	2.07	a
	[0.939, 0.950]	[18.98, 19.05]				
BSR43	0.837 $\pm$ 0.002	18.43 $\pm$ 0.02	2.91	4.80	2.19	b
	[0.833, 0.842]	[18.39, 18.46]				
BSR75	0.688 $\pm$ 0.002	18.33 $\pm$ 0.02	3.59	6.30	2.02	c
	[0.684, 0.691]	[18.29, 18.36]				
SSR20	0.726 $\pm$ 0.003	29.02 $\pm$ 0.04	3.98	7.08	2.18	d
	[0.719, 0.733]	[28.94, 29.10]				
SSR30	0.658 $\pm$ 0.003	28.66 $\pm$ 0.04	4.39	8.47	2.30	e
	[0.652, 0.665]	[28.58, 28.75]				
SSR73	0.539 $\pm$ 0.003	26.93 $\pm$ 0.05	4.47	10.57	2.29	f
	[0.534, 0.544]	[26.84, 27.01]				

^a R_N = number of fibers ≤ 0.1 mm divided by number of fibers > 0.1 mm.

^b Width skewness computed for fibers with length > 0.1 mm only, to reduce the influence of abundant short fragments.

^c Games–Howell post hoc grouping on L_1 ($\alpha = 0.05$), samples sharing a letter are not significantly different. All pairwise comparisons are significant, hence each sample receives a unique letter.

margin could be measured against known ground truth. On the manually annotated real junctions the same comparison is much closer, and we do not claim the advantage to be large there, which we attribute largely to the heuristic rule having been tuned on industrial samples of this kind. This highlights that optimal feature combination—not merely feature selection—drives the performance improvement.

A natural question for any data-driven method trained exclusively on synthetic data is whether the learned model transfers to real images. The agreement with the FS5 analyzer on industrial samples, and the consistent results obtained on bamboo and softwood pulps without retraining, provide evidence of transferability. Several factors may underpin it. The four geometric features occupy similar numerical ranges for synthetic and real intersections, which makes the learned scoring function tolerant of the photometric difference between the two domains. The synthetic generator, although agnostic to specific pulp species, reproduces the geometric phenomena that dominate real intersections—variable crossing angles, width heterogeneity and curvature near junctions. A further qualification concerns width. Backlight illumination yields a silhouette, so the measured quantity is the projected width of the fiber; because pulp fibers are flattened ribbons, this projected width varies with the twist of the fiber about its own axis. This is a residual source of uncertainty, and it is one reason why we do not present the width measurement as exact. We do not regard the consistency observed on the cross-material samples as proof that no systematic measurement error is present, since no independent reference was available for those samples.

Beyond measurement accuracy—a mean relative

deviation of 2.13% across all metrics—the method offers several operational advantages that are relevant to industrial deployment. The statistical sensitivity demonstrated in our cross-material experiment ($\eta^2 \geq 0.99$, all pairwise comparisons significant) confirms that the pipeline can reliably distinguish process states that differ only modestly in fiber morphology, a capability directly relevant to quality control and refining optimization. Moreover, the capacity to resolve the full fiber population distribution—including the heavy concentration of short fragments generated by refining—provides process-sensitive descriptors such as the short-to-long fiber ratio R_N and distribution skewness, which are unavailable from bulk optical analyzers. The computational efficiency of the pipeline is compatible with laboratory quality-control throughput, and the software-based implementation can be deployed on standard workstations without specialized hardware, offering flexibility that dedicated instruments cannot easily provide.

CONCLUSION AND FUTURE WORK

This paper proposes a data-driven framework for identifying and resolving fiber intersections in pulp and paper micrographs by replacing conventional heuristic rules with KAN-based multi-feature aggregation. At each candidate breakpoint, four local geometric features—direction angle, local width, curvature, and cross-region relation—are extracted and combined by the KAN to score fragment pairs, with final matching performed by a greedy algorithm. On the synthetic benchmark with ground truth, the method attains an

F1 score of 0.978 at a single fiber density (15 curves per image); averaged over the 15 tested densities it reaches 0.9656. Applied to industrial pulp images, it yields length and width distributions that closely match a commercial FS5 analyzer while avoiding the junction-related biases of the heuristic matching baseline. These results establish KAN-based feature aggregation as a quantitatively validated alternative to heuristic matching in automated fiber morphology characterization for cellulose-based materials: the learned matching reaches an F1 of 0.94 against the manual annotation of real junctions, the measurements are highly reproducible, and the approach resolves the junction-related bias of the heuristic rule at a computational cost low enough for throughput-oriented use, although the margin over that rule is small on real junctions and the absence of systematic error remains unverified for the cross-material samples. Despite these promising results, the method still has limitations. It is based on 2D micrographs, so projection errors from overlapping fibers cannot be fully avoided. In addition, the manual annotation of real junctions reported above verifies that the full four-feature model ranks highest on real junctions, but the margin over both the reduced configurations and the heuristic baseline is narrow on that annotation (F1 0.9382 versus 0.9144 for the heuristic baseline), so no claim of a large matching advantage on real data is warranted. No independent instrument reference is available for the six cross-material samples, so the absence of systematic measurement error cannot be established for them.

Several directions warrant further investigation. First, extending the framework to 3D imaging—e.g., confocal z-stacks or X-ray microtomography—would mitigate projection artifacts; the learned matching concept transfers naturally using 3D geometric features. Second, the current KAN operates locally on each intersection region; incorporating global context via graph neural networks could resolve ambiguities that are locally underdetermined. Third, a semi-synthetic dataset will be built by extracting real fibers and placing them randomly within synthetic images, so that training and testing can be performed on data that combines real fiber geometry with known ground truth. Fourth, the direction and cross-region angles are treated as scalar features (Section "Breakpoint Feature Extraction"); representing them on the circle, for example by sine/cosine pairs or as explicit angular differences, would remove the $0^\circ/360^\circ$ discontinuity and allow higher-dimensional angular descriptors to be introduced, and a systematic comparison of circular encodings is therefore a natural extension of the present feature set. Finally, online learning from operator feedback could enable progressive adaptation to individual mill conditions, offering self-improving

quality-control capabilities for the pulp and paper industry.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

FUNDING

This work was supported by the National Natural Science Foundation of China (Grant No.: 92467106, 62273151, W2521175), the Guangdong Generic Institution Innovation Team Research Foundation (Grant No.: 2023KCXTDO72).

CONFLICTS OF INTEREST

The authors declare no conflicts of interest.

REFERENCES

- Alejo AM, Beltrán CA (2024). A deep learning approach to distance map generation applied to automatic fiber diameter computation from digital micrographs. *Sensors* 24:5497.
- Arafat Y, Cuesta-Apausa C, Castellano E, Reyes-Aldasoro CC (2025). Fibre tracing in biomedical images: an objective comparison between seven algorithms. *PLoS ONE* 20:e0320006.
- Baez C, Ringania U, Bhamla S, Moon RJ (2026). Semi-automated image analysis of cellulose nanofibrils using machine learning segmentation and morphological thinning. *Cellulose* 33:3259–82.
- Blum H (1967). A transformation for extracting new descriptors of shape. In: Wathen-Dunn W, ed., *Models for the perception of speech and visual form*. Cambridge: MIT Press, 362–80.
- Bradley D, Roth G (2007). Adaptive thresholding using the integral image. *J Graph Tools* 12:13–21.
- Chen W, Liu M, Zhan Q, Tan Y, Meijering E, Radojevic M, Wang Y (2021). Spherical-patches extraction for deep-learning-based critical points detection in 3D neuron microscopy images. *IEEE Trans Med Imaging* 40:527–38.
- Chirakitsakul T, Wattuya P, Somboon P, Jansakr P, Watcharopas C (2025). Integration of convolutional neural network and image processing for pulp fibril detection and measurement. *IEEE Access* 13:74630–42.

- D'Amore A, Stella JA, Wagner WR, Sacks MS (2010). Characterization of the complete fiber network topology of planar fibrous tissues and scaffolds. *Biomaterials* 31:5345–54.
- Deng Z, Yu L, Wang L, Ke W (2024). An algorithm for cross-fiber separation in yarn hairiness image processing. *Vis Comput* 40:3591–9.
- Dias PAN, Rodrigues R, Reis MS (2024). Experimental assessment of paper formation conditions and structural two-sidedness and their impacts on curl phenomena. *Processes* 12:1536.
- Gaiselmann G, Manke I, Lehnert W, Schmidt V (2013). Extraction of curved fibers from 3D data. *Image Anal Stereol* 32:57–63.
- Guo C, Liu M, Guan T, Chen W, Wen H, Zeng T, Wang Y (2021). Crossover structure separation with application to neuron tracing in volumetric images. *IEEE Trans Instrum Meas* 70:1–13.
- Howard CV, Reed MG (2005). *Unbiased Stereology*. Abingdon: Garland Science/BIOS Scientific, 2nd ed.
- Huang X, Wen D, Zhao Y, Wang Q, Zhou W, Deng D (2016). Skeleton-based tracing of curved fibers from 3D X-ray microtomographic imaging. *Results Phys* 6:170–7.
- Kayali Y, Gleadall A, Silberschmidt VV (2025). Geometric characterisation of stochastic fibrous networks: a comprehensive review. *Fibers* 13:27.
- Lee TC, Kashyap RL, Chu CN (1994). Building skeleton models via 3-D medial surface axis thinning algorithms. *CVGIP Graph Models Image Process* 56:462–78.
- Lin J, Feng L, Yin P, Liu Y (2026). Achieving synergistic morphology–topology analysis of microscale fiber networks in pulp and paper via a collaborative computational framework. In: Twelfth International Conference on Mechanical Engineering, Materials and Automation Technology (MMEAT 2026), vol. 14355 of *Proc. SPIE*. SPIE.
- Lindström SB, Amjad R, Gåhlin E, Andersson L, Kaarto M, Liubytka K, Persson J, Berg JE, Engberg BA, Nilsson F (2024). Pulp particle classification based on optical fiber analysis and machine learning techniques. *Fibers* 12:2.
- Liu Z, Wang Y, Vaidya S, Ruehle F, Halverson J, Soljačić M, Hou TY, Tegmark M (2024). KAN: Kolmogorov-Arnold networks. In: International Conference on Learning Representations (ICLR 2025). OpenReview.
- Lux J (2013). Automatic segmentation and structural characterization of low density fibreboards. *Image Anal Stereol* 32:13–25.
- Muñoz A, Docaj A, Fernandez J, Carriero A (2025). FiberO for an automated quantitative analysis of fibers orientation and organization in biological fibrous tissues. *Front Bioeng Biotechnol* 12:1497837.
- Otsu N (1979). A threshold selection method from gray-level histograms. *IEEE Trans Syst Man Cybern* 9:62–6.
- Paunonen S, Mattsson A, Pettersson G, Ketoja JA (2026). Property variations of binder-free lignin-rich fiber networks driven by forming processes and hot pressing. *ACS Omega* 11:7129–41.
- Pennells J, Heuberger B, Chaléat C, Martin DJ (2022). Assessing cellulose micro/nanofibre morphology using a high throughput fibre analysis device to predict nanopaper performance. *Cellulose* 29:2599–616.
- Pourdeyhimi B, Dent R (1999). Measuring fiber diameter distribution in nonwovens. *Text Res J* 69:233–6.
- Qamar S, Baba AI, Verger S, Andersson M (2024). Segmentation and characterization of macerated fibers and vessels using deep learning. *Plant Methods* 20:126.
- Ronneberger O, Fischer P, Brox T (2015). U-Net: convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, vol. 9351 of *Lecture Notes in Computer Science*. Berlin: Springer. 234–41.
- Soille P (2003). *Morphological image analysis: principles and applications*. Berlin: Springer, 2nd ed.
- Sten O, Del Dottore E, Pugno N, Mazzolai B (2024). A ridge-based detection algorithm with filament overlap identification for 2D mycelium network analysis. *Ecol Inform* 82:102670.
- Togo R, Cardona S, Nagle I, Koenderink GH, Fereidoonhezahad B, Peirlinck M (2026). ToFiE, a topology-aware fiber extraction workflow for 3D reconstruction of dense and heterogeneous biological fiber networks from microscopy images. *arXiv preprint arXiv:2604.18230*.
- Underwood EE (1970). *Quantitative Stereology*. Reading, MA: Addison-Wesley.
- Verho T, Turpeinen T, Asad F, Immonen K (2025). A skeletonization-based approach for individual fiber separation in tomography images of biocomposites. *Comput Mater Sci* 246:113372.